

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
27 May 2004 (27.05.2004)

PCT

(10) International Publication Number
WO 2004/045162 A2

(51) International Patent Classification: **H04L 12/56**

(21) International Application Number:
PCT/GB2003/004893

(22) International Filing Date:
11 November 2003 (11.11.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0226249.1 11 November 2002 (11.11.2002) GB

(71) Applicant (for all designated States except US): **CLEAR-SPEED TECHNOLOGY LIMITED** [GB/GB]; 3110 Great Western Court, Hunts Ground Road, Stoke Gifford, Bristol BS34 8HP (GB).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **SPENCER, Anthony** [GB/GB]; 34 Amberley Way, Wickwar, Wotton-under-Edge, South Gloucestershire GL12 8LP (GB).

(74) Agents: **O'CONNELL, David, Christopher et al.**; Haseltine Lake, Imperial House, 15-19 Kingsway, London WC2B 6UD (GB).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (utility model), DE, DK (utility model), DK, DM, DZ, EC, EE (utility model), EE, EG, ES, FI (utility model), FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.


(84) Designated States (*regional*): ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

[Continued on next page]

(54) Title: **TRAFFIC MANAGEMENT ARCHITECTURE**



(57) Abstract: An architecture for sorting incoming data packets in real time, on the fly, processes the packets and places them into an exit order queue before storing the packets. This is in contrast to the traditional way of storing first then sorting later and provides rapid processing capability. A processor (22) generates packet records from an input stream (20) and determines an exit order number for the related packet. The records are stored in an orderlist manager (24) whilst the data portions are stored in memory hub (21) for later retrieval in the exit order stored in the manager (24). The processor (22) is preferably a parallel processor array using SIMD and provided with rapid access to shared state (23) by a state engine.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

TRAFFIC MANAGEMENT ARCHITECTURE

Field of the Invention

The present invention concerns the management of traffic, such as data and
5 communications traffic, and provides an architecture for a traffic manager that surpasses known traffic management schemes in terms of speed, efficiency and reliability.

Background to the Invention

The problem that modern traffic management schemes have to contend with is the sheer volume. Data arrives at a traffic handler from multiple sources at unknown rates and
10 volumes and has to be received, sorted and passed on "on the fly" to the next items of handling downstream. Received data may be associated with a number of attributes by which priority allocation, for example, is applied to individual data packets or streams, depending on the class of service offered to an individual client. Some traffic may therefore have to be queued whilst later arriving but higher priority traffic is processed.
15 A router's switch fabric can deliver packets from multiple ingress ports to one of a number of egress ports. The linecard connected to this egress port must then transmit these packets over some communication medium to the next router in the network. The rate of transmission is normally limited to a standard rate. For instance, an OC-768 link would transmit packets over an optical fibre at a rate of 40 Gbits/s.
20 With many independent ingress paths delivering packets for transmission at egress, the time-averaged rate of delivery cannot exceed 40 Gbits/s for this example. Although over time the input and output rates are equivalent, the short term delivery of traffic by the fabric is "bursty" in nature with rates often peaking above the 40 Gbits/s threshold. Since the rate of receipt can be greater than the rate of transmission, short term packet queueing
25 is required at egress to prevent packet loss. A simple FIFO queue is adequate for this purpose for routers which provide a flat grade of service to all packets. However, more complex schemes are required in routers which provide Traffic Management. In a converged internetwork, different end user applications require different grades of service in order to run effectively. Email can be carried on a best effort service where no
30 guarantees are made regarding rate of or delay in delivery. Real-time voice data has a much more demanding requirement for reserved transmission bandwidth and guaranteed minimum delay in delivery. This cannot be achieved if all traffic is buffered in the same

FIFO queue. A queue per so-called "Class of Service" is required so that traffic routed through higher priority queues can bypass that in lower priority queues. Certain queues may also be assured a guaranteed portion of the available output line bandwidth.

On first sight the traffic handling task appears to be straightforward. Packets are placed in queues according to their required class of service. For every forwarding treatment that a system provides, a queue must be implemented. These queues are then managed by the following mechanisms:

- Queue management assigns buffer space to queues and prevents overflow
- Measures are implemented to cause traffic sources to slow their transmission rates if queues become backlogged
- Scheduling controls the de-queuing process by dividing the available output line bandwidth between the queues.

Different service levels can be provided by weighting the amount of bandwidth and buffer space allocated to different queues, and by prioritised packet dropping in times of congestion. Weighted Fair Queueing (WFQ), Deficit Round Robin (DRR) scheduling, Weighted Random Early Detect (WRED) are just a few of the many algorithms which might be employed to perform these scheduling and congestion avoidance tasks.

In reality, system realisation is confounded by some difficult implementation issues:

- High line speeds can cause large packet backlogs to rapidly develop during brief congestion events. Large memories of the order 500 MBytes to 1 GBytes are required for 40 Gbits/s line rates.
- The packet arrival rate can be very high due to overspeed in the packet delivery from the switch fabric. This demands high data read and write bandwidth into memory. More importantly, high address bandwidth is also required.
- The processing overhead of some scheduling and congestion avoidance algorithms is high.

Priority queue ordering for some (FQ) scheduling algorithms is a non-trivial problem at high speeds.

- A considerable volume of state must be maintained in support of scheduling and congestion avoidance algorithms, to which low latency access is required. The volume of state increases with the number of queues implemented.

- As new standards and algorithms emerge, the specification is a moving target. To find a flexible (ideally programmable) solution is therefore a high priority.

In a conventional approach to traffic scheduling, one might typically place packets directly into an appropriate queue on arrival, and then subsequently dequeue packets from those queues into an output stream.

Figure 1 shows the basic layout of the current approach to traffic management. It can be thought of as a “queue first, think later” strategy. Data received at the input is split into a number of queues in parallel channels 2.1 to 2.n. A traffic scheduler processor 3 receives the data from the parallel channels and sorts them into order. The order may be determined by the priority attributes, for example, mentioned above. State is stored in memory 4 accessible by the processor. The output from the processor represents the new order as determined by the processor in dependence on the quality of service attributes assigned to the data at the outset.

The traffic scheduler 3 determines the order of de-queuing. Since the scheduling decision can be processing-intensive as the number of input queues increases, queues are often arranged into small groups which are locally scheduled into an intermediate output queue. This output queue is then the input queue to a following scheduling stage. The scheduling problem is thus simplified using a “divide-and-conquer” approach, whereby high performance can be achieved through parallelism between groups of queues in a tree type structure, or so-called hierarchical link sharing scheme.

This approach works in hardware up to a point. For the exceptionally large numbers of input queues (of the order 64k) required for per-flow traffic handling, the first stage becomes unmanageably wide to a point that it becomes impractical to implement the required number of schedulers.

Alternatively, in systems which aggregate all traffic into a small number of queues parallelism between hardware schedulers cannot be exploited. It then becomes extremely difficult to implement a single scheduler - even in optimised hardware - that can meet the required performance point.

With other congestion avoidance and queue management tasks to perform in addition to scheduling, it is apparent that a new approach to traffic handling is required. The queue first, think later strategy often fails and data simply has to be jettisoned. There is

therefore a need for an approach to traffic management that does not suffer from the same defects as the prior art and does not introduce its own fallibilities.

Summary of the invention

In one aspect, the invention provides a system comprising means for sorting incoming
5 data packets in real time before said packets are stored in memory.

In another aspect, the invention provides a data packet handling system, comprising means whereby incoming data packets are assigned an exit order before being stored in memory.

In yet another aspect, the invention provides a method for sorting incoming data packets
10 in real time, comprising sorting the packets into an exit order before storing them in memory.

The sorting means may be responsive to information contained within a packet and/or within a table and/or information associated with a data packet stream in which said packet is located, whereby to determine an exit order number for that packet.

15 The packets may be inserted into one or more queues by a queue manager adapted to insert packets into the queue means in exit order. There may be means to drop certain packets before being output from said queue means or before being queued in the queue means.

The system may be such that the sorting means and the queue means process only packet
20 records containing information about the packets, whereas data portions of the packets are stored in the memory for output in accordance with an exit order determined for the corresponding packet record.

The sorting means preferably comprises a parallel processor, such as an array processor, more preferably a SIMD processor.

25 There may be further means to provide access for the parallel processors to shared state. A state engine may control access to the shared state.

Tables of information for sorting said packets or said packet records may be provided, wherein said tables are stored locally to each processor or to each processor element of a parallel processor. The tables may be the same on each processor or on each processor
30 element of a parallel processor. The tables may be different on different processors or on different processor elements of a parallel processor.

The processors or processor elements may share information from their respective tables,

such that: (a) the information held in the table for one processor is directly accessible by a different processor or the information held in the table in one processor element may be accessible by other processing element(s) of the processor; and (b) processors may have access to tables in other processors or processor elements have access to other processor elements in the processor, whereby processors or processor elements can perform table lookups on behalf of other processor(s) or processor elements of the processor.

The invention also encompasses a computer system, comprising a data handling system as previously specified; a network processing system, comprising a data handling system as previously specified; and a data carrier containing program means adapted to perform a corresponding method.

Brief Description of the Drawings

The invention will be described with reference to the following drawings, in which:

Figure 1 is a schematic representation of a prior art traffic handler; and

Figure 2 is a schematic representation of a traffic handler in accordance with the invention.

Detailed Description of the Illustrated Embodiments

The present invention turns current thinking on its head. Figure 2 shows schematically the basic structure underlying the new strategy for effective traffic management. It could be described as a “think first, queue later TM” strategy.

Packet data (traffic) received at the input 20 has the header portions stripped off and record portions of fixed length generated therefrom, containing information about the data, so that the record portions and the data portions can be handled separately. Thus, the data portions take the lower path and are stored in Memory Hub 21. At this stage, no attempt is made to organise the data portions in any particular order. However, the record portions are passed to a processor 22, such as a SIMD parallel processor, comprising one or more arrays of processor elements (PEs). Typically, each PE contains its own processor unit, local memory and register(s).

In contrast to the prior architecture outlined in Figure 1, the present architecture shares state 23 in the PE arrays under the control of a State Engine (not shown) communicating with the PE array(s). It should be emphasised that only the record portions are processed in the PE array. The record portions are all the same length, so their handling is predictable, at least in terms of length.

The record portions are handled in the processor 22. Here, information about the incoming packets is distributed amongst the PEs in the array. This array basically performs the same function as the processor 3 in the prior art (Figure 1) but the operations are spread over the PE array for vastly more rapid processing. This processing effectively
5 “time-stamps” the packet records to indicate when the corresponding data should be exited, assuming that it should actually be exited and not jettisoned, for example.

The results of this processing are sent to the orderlist manager 24, which is an “intelligent” queue system which places the record portions in the appropriate exit order, for example in bins allocated to groups of data exit order numbers. The manager 24 is
10 preferably dynamic, so that new data packets with exit numbers having a higher priority than those already in an appropriate exit number bin can take over the position previously allocated. It should be noted that the PE array 22 simply calculates the order in which the data portions are to be output but the record portions themselves do not have to be put in that order. In other words, the PEs do not have to maintain the order of packets being
15 processed nor sort them before they are queued.

Previous systems in which header and data portions were treated as one entity became unwieldy, slow and cumbersome because of the innate difficulty of preserving the integrity of the whole packet yet still providing enough bandwidth to handle the combination. In the present invention, it is only necessary for the Memory Hub 21 to
20 provide sufficient bandwidth to handle just the data portions. The memory hub can handle packets streaming in at real time. The memory hub can nevertheless divide larger data portions into fragments, if necessary, and store them in physically different locations, provided, of course, there are pointers to the different fragments to ensure read out of the entire content of such data packets.

25 In order to overcome the problem of sharing state over all the PEs in the array, multiple PEs are permitted to access (and modify) the state variables. Such access is under the control of a State Engine (not shown), which automatically handles the “serialisation” problem of parallel access to shared state.

The output 25, in dependence on the exit order queue held in the Orderlist Manager 24,
30 instructs the Memory Hub 21 to read out the corresponding packets in that required order, thereby releasing memory locations for newly received data packets in the process.

The chain-dotted line 26 enclosing the PE array 22, shared state/State Engine 23 and Orderlist Manager 24 signifies that this combination of elements can be placed on a single chip and that this chip can be replicated, so that there may be one or two (or more) chips interfacing with single input 20, output 25 and Memory Hub 21. As is customary, the chip will also include necessary additional components, such as a distributor and a collector per PE array to distribute data to the individual PEs and to collect processed data from the PEs, plus semaphore block(s) and interface elements.

The following features are significant to the new architecture:

- There are no separate, physical stage one input queues.
- Packets are effectively sorted directly into the output queue on arrival. A group of input queues thus exists in the sense of being interleaved together within the single output queue.

- These interleaved “input queues” are represented by state in the queue state engine.

This state may track queue occupancy, finish time/number of the last packet in the queue

etc. Occupancy can be used to determine whether or not a newly arrived packet should be placed in the output queue or whether it should be dropped (congestion management).

Finish numbers are used to preserve the order of the “input queues” within the output queue and determine an appropriate position in the output queue for newly arrived packets (scheduling).

- Scheduling and congestion avoidance decisions are thus made “on the fly” prior to enqueueing (ie “Think first, queue later”TM).

- This technique is made possible by the deployment of a high performance data flow processor which can perform the required functions at wire speed. Applicant’s array processor is ideal for this purpose, providing a large number of processing cycles per packet for packets arriving at rates as high as one every couple of system clock cycles.

Ancillary features

Class of Service (CoS) tables:

CoS parameters are used in scheduling and congestion avoidance calculations. They are conventionally read by processors as a fixed group of values from a class of service table in a shared memory. This places further demands on system bus and memory access bandwidth. The table size also limits the number of different classes of service which may be stored.

An intrinsic capability of Applicant's array processor is rapid, parallel local memory access. This can be used to advantage as follows:

- The Class of Service table is mapped into each PE's memory. This means that all passive state does not require lookup from external memory. The enormous internal memory addressing bandwidth of SIMD processor is utilised.

- By performing multiple lookups into local memories in a massively parallel fashion instead of single large lookups from a shared external table there is a huge number of different Class of Service combinations available from a relatively small volume of memory.

Table sharing between PEs - PEs can perform proxy lookups on behalf of each other. A single CoS table can therefore be split across two PEs, thus halving the memory requirement.

Summary

It can thus be appreciated that the present invention is capable of providing the following key features, marking considerable improvements over the prior art:

- Traditional packet scheduling involves parallel enqueueing and then serialised scheduling from those queues. For high performance traffic handling we have turned this around. Arriving packets are first processed in parallel and subsequently enqueued in a serial orderlist. This is referred to as "Think First Queue Later"TM

- The deployment of a single pipeline parallel processing architecture (Applicant's array processor) is innovative in a Traffic Handling application. It provides the wire speed processing capability which is essential for the implementation of this concept.

- An alternate form of parallelism (compared to independent parallel schedulers) is thus exploited in order to solve the processing issues in high speed Traffic Handling.

Claims

1. A system comprising means for sorting incoming data packets in real time before said packets are stored in memory.
- 5 2. A data packet handling system, comprising means whereby incoming data packets are assigned an exit order before being stored in memory.
- 10 3. A system as claimed in claim 1 or claim 2 wherein the sorting means is responsive to information contained within a packet whereby to determine an exit order number for that packet.
- 15 4. A system as claimed in claim 2, wherein the sorting means is responsive to information contained in a table whereby to determine an exit order number for that packet.
- 20 5. A system as claimed in claim 2, wherein the sorting means is responsive to information associated with a data packet stream in which said packet is located whereby to determine an exit order number for that packet.
- 25 6. A system as claimed in claim 1 or claim 2, comprising queue means to queue sorted packets for output in exit order.
7. A system as claimed in claim 6, wherein said sorting means is adapted to insert sorted packets in said queue means in exit order.
8. A system as claimed in claim 6 or 7, wherein said queue means is a single queue.
- 30 9. A system as claimed in claim 8, wherein said single queue provides a plurality of virtual queues.

10. A system as claimed in claim 6, further comprising a queue manager adapted to insert packets into said queue means in exit order.
- 5 11. A system as claimed in claim 6, further comprising means to drop certain packets before being output from said queue means.
12. A system as claimed in claim 6, further comprising means to drop certain packets before being queued in said queue means.
- 10 13. A system as claimed in any of the preceding claims, wherein:
said sorting means and said queue means process only packet records
containing information about said packets, and
data portions of said packets are stored in said memory for output in
accordance with an exit order determined for the corresponding
15 packet record.
14. A system as claimed in any of the preceding claims, wherein said sorting means comprises a parallel processor.
- 20 15. A system as claimed in claim 14, wherein said parallel processor is an array processor.
16. A system as claimed in claim 14, wherein said array processor is a SIMD
processor.
- 25 17. A system as claimed in claim 14, 15 or 16, further comprising means to provide access for said parallel processors to shared state.
18. A system as claimed in claim 17, further comprising a state engine to control
30 said access to said shared state.

19. A system as claimed in any of claims 1 to 18, further comprising tables of information for sorting said packets or said packet records, wherein said tables are stored locally to each processor or to each processor element of a parallel processor.

5

20. A system as claimed in claim 19, wherein said tables are the same on each processor or on each processor element of a parallel processor.

21. A system as claimed in claim 19, wherein said tables are different on different processors or on different processor elements of a parallel processor.

10

22. A system as claimed in claim 19, wherein said processors or processor elements share information from their respective tables, such that:

(a) the information held in the table for one processor is directly accessible by a different processor or the information held in the table in one processor element is accessible by other processing element(s) of the processor; and

15

(b) processors have access to tables in other processors or processor elements have access to other processor elements in the processor, whereby processors or processor elements can perform table lookups on behalf of other processor(s) or processor elements of the processor.

20

23. A system as claimed in any of the preceding claims, wherein said sorting means implement algorithms for packet scheduling in accordance with predetermined criteria, such as WFQ, DFR, congestion avoidance (eg WRED) or other prioritisation and sorting.

25

24. A method for sorting incoming data packets in real time, comprising sorting the packets into an exit order before storing them in memory.

30

25. A method as claimed in claim 24, wherein the sorting is responsive to information contained within a packet whereby to assign an exit order number for that packet.
- 5 26. A method as claimed in claim 24, wherein the sorting is responsive to information contained in a table whereby to determine an exit order number for that packet.
- 10 27. A method as claimed in claim 24, wherein the sorting is responsive to information associated with a data packet stream in which said packet is located whereby to determine an exit order number for that packet.
28. A method as claimed in claim 24, further comprising queuing sorted packets for output in exit order.
- 15 29. A method as claimed in claim 28, wherein said packets are inserted into a queue means in exit order determined by the means performing the sorting.
- 20 30. A method as claimed in claim 28, comprising inserting sorted packets into a queue means in exit order under control of a queue manager.
31. A method as claimed in claim 29 or 30, wherein said queuing is performed using a single output queue.
- 25 32. A method as claimed in claim 31, further comprising providing a plurality of virtual queues by means of said single output queue.
33. A method as claimed in claim 28, further comprising dropping certain packets before being output from said queue means.
- 30 34. A method as claimed in claim 28, further comprising dropping certain packets before being queued in said queue means.

35. A method as claimed in any of claims 24-34, wherein:

said sorting and said queuing operations are performed only on packet records containing information about said packets, said method further comprising:

storing data portions of said packets in said memory for output in accordance with an exit order number determined for the corresponding packet record.

36. A method as claimed in any of claims 24-34, wherein said sorting is performed by a parallel processor.

37. A method as claimed in claim 36, wherein said parallel processor is an array processor.

38. A method as claimed in claim 36, wherein said array processor is a SIMD processor.

39. A method as claimed in claim 36, 37 or 38, further comprising providing access for said processors to shared state under control of a state engine.

40. A method as claimed in claim 39, further comprising providing tables of information for sorting said packets or said packet records, wherein said tables are stored locally to each processor or to each processor element of a parallel processor.

41. A method as claimed in claim 40, wherein said tables are the same on each processor or on each processor element of a parallel processor.

42. A method as claimed in claim 40, wherein said tables are different on different processors or on each processor element of a parallel processor.

43. A method as claimed in claim 40, wherein said processors or processor elements share information from their respective tables, such that:

(a) the information held in the table for one processor is made directly accessible by a different processor or the information held in the table of one processor element is made directly accessible to other processor element(s) of the processor; and

(b) access is provided for said processor or processor elements to tables in other processors or processor elements, whereby processors or processor elements can perform table lookups on behalf of another processor or processor element.

44. A system as claimed in any of claims 1-23, wherein said sorting means implement algorithms for packet scheduling in accordance with predetermined criteria, such as WFQ, DFR, congestion avoidance (eg WRED) or other prioritisation and sorting.

45. A computer system, comprising a data handling system as claimed in any of claims 1-23.

46. A network processing system, comprising a data handling system as claimed in any of claims 1-23.

47. A computer system adapted to perform the method as claimed in any of claims 24-43.

48. A network processing system adapted to perform the method as claimed in any of claims 24-43.

49. A computer system as claimed in claim 45 implemented as one or more silicon integrated circuits.

50. A data carrier containing program means adapted to perform the method as claimed in any of claim 24 to 43.

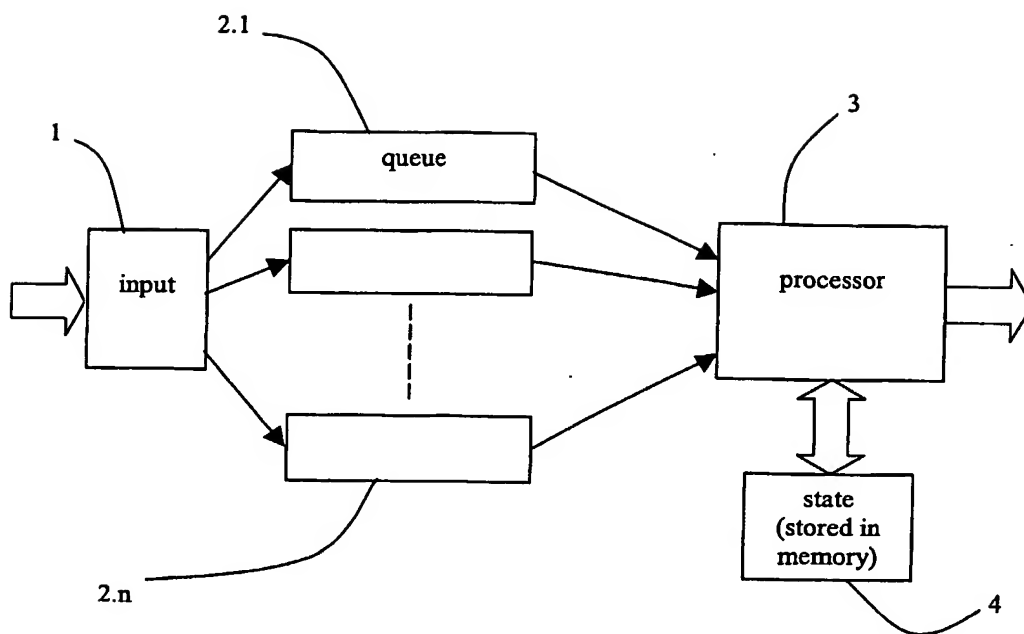


Fig 1




Fig 2

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
27 May 2004 (27.05.2004)

PCT

(10) International Publication Number
WO 2004/045162 A3

(51) International Patent Classification⁷: **H04L 12/56**

(21) International Application Number:
PCT/GB2003/004893

(22) International Filing Date:
11 November 2003 (11.11.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
0226249.1 11 November 2002 (11.11.2002) GB

(71) Applicant (for all designated States except US): **CLEAR-SPEED TECHNOLOGY LIMITED** [GB/GB]; 3110 Great Western Court, Hunts Ground Road, Stoke Gifford, Bristol BS34 8HP (GB).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **SPENCER, Anthony** [GB/GB]; 34 Amberley Way, Wickwar, Wotton-under-Edge, South Gloucestershire GL12 8LP (GB).


(74) Agents: **O'CONNELL, David, Christopher et al.**; Haseltine Lake, Imperial House, 15-19 Kingsway, London WC2B 6UD (GB).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (utility model), DE, DK (utility model), DK, DM, DZ, EC, EE (utility model), EE, EG, ES, FI (utility model), FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (*regional*): ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: **TRAFFIC MANAGEMENT ARCHITECTURE**



(57) Abstract: An architecture for sorting incoming data packets in real time, on the fly, processes the packets and places them into an exit order queue before storing the packets. This is in contrast to the traditional way of storing first then sorting later and provides rapid processing capability. A processor (22) generates packet records from an input stream (20) and determines an exit order number for the related packet. The records are stored in an orderlist manager (24) whilst the data portions are stored in memory hub (21) for later retrieval in the exit order stored in the manager (24). The processor (22) is preferably a parallel processor array using SIMD and provided with rapid access to shared state (23) by a state engine.



Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(88) Date of publication of the international search report:
16 September 2004

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 03/04893

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the International search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2002/075882 A1 (DATTA UTPAL ET AL) 20 June 2002 (2002-06-20) paragraph '0033! paragraph '0102! -----	1-13, 43-50
X	EP 1 137 225 A (CIT ALCATEL) 26 September 2001 (2001-09-26) paragraph '0022! - paragraph '0028! -----	1-50
X	EP 0 372 795 A (AMERICAN TELEPHONE & TELEGRAPH) 13 June 1990 (1990-06-13) column 3, paragraph 37 - column 4, paragraph 23 -----	1,2

☐ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the International search 20 July 2004	Date of mailing of the international search report 26/07/2004
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Authorized officer Gregori, S

INTERNATIONAL SEARCH REPORT

information on patent family members

International Application No

PCT/GB 03/04893

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 2002075882	A1	20-06-2002	AU 3883499 A	23-11-1999
			CA 2331820 A1	11-11-1999
			EP 1080563 A1	07-03-2001
			WO 9957858 A1	11-11-1999
EP 1137225	A	26-09-2001	EP 1137225 A1	26-09-2001
			JP 2001285330 A	12-10-2001
			US 2001030974 A1	18-10-2001
EP 0372795	A	13-06-1990	US 4914650 A	03-04-1990
			AU 609341 B2	26-04-1991
			AU 4560689 A	14-06-1990
			CA 2000557 A1	06-06-1990
			EP 0372795 A2	13-06-1990
			JP 2202250 A	10-08-1990